

---

# Stellar Spectra Fitting with Amortized Neural Posterior Estimation and *nbi*

---

Keming Zhang<sup>1</sup> Tharindu Jayasinghe<sup>1</sup> Joshua Bloom<sup>1</sup>

## Abstract

Modern surveys often deliver hundreds of thousands of stellar spectra at once, which are fit to spectral models to derive stellar parameters/labels. Therefore, the technique of Amortized Neural Posterior Estimation (ANPE) stands out as a suitable approach, which enables the inference of large number of targets as sub-linear/constant computational costs. Leveraging our new *nbi* software package<sup>1</sup>, we train an ANPE model for the APOGEE survey and demonstrate its efficacy on both mock and real APOGEE stellar spectra. Unique to the *nbi* package is its out-of-the-box functionality on astronomical inverse problems with sequential data. As such, we have been able to acquire the trained model with minimal effort. We introduce an effective approach to handling the measurement noise properties inherent in spectral data, which utilizes the actual uncertainties in the observed data. This allows training data to resemble observed data, an aspect that is crucial for ANPE applications. Given the association of spectral data properties with the observing instrument, we discuss the utility of an ANPE “model zoo,” where models are trained for specific instruments and distributed under the *nbi* framework to facilitate real-time stellar parameter inference.

## 1. Introduction

The efficacy of Amortized Neural Posterior Estimation (ANPE) has been demonstrated for a wide range of astronomical problems, such as gravitational waves (Dax et al., 2021), galaxy SED fitting (Hahn & Melchior, 2022; Khullar et al., 2022), binary-lens microlensing (Zhang et al., 2021),

<sup>1</sup>Department of Astronomy, University of California, Berkeley, Berkeley, USA. Correspondence to: Keming Zhang <kemingz@berkeley.edu>.

*ICML 2023 Workshop on Machine Learning for Astrophysics*, Honolulu, Hawaii, USA. PMLR 202, 2023. Copyright 2023 by the author(s).

<sup>1</sup>Accepted as a separate paper to this workshop: “*nbi*: the Astronomer’s Package for Neural Posterior Estimation.” Code at: <https://github.com/kmzzhang/nbi>

exoplanet atmosphere retrieval (Vasist et al., 2023), among others. The predicate of ANPE is that the observed data are drawn from the same distribution as the training data. This requirement has substantially complicated the application of ANPE to time-domain astronomy. For light-curve data extracted from the same time-domain survey, each photometric data point is taken under different conditions, and the light-curve time samplings can vastly differ, with large temporal gaps, some sampled very sparsely, others densely.

In comparison, spectroscopy surveys are much better positioned in the application of ANPE. Compared to light-curve datasets, all data points for any given spectra are taken at the same time under the same observing condition and with the same instrument. For any given spectrograph, the wavelength grid/sampling has fixed options, and the noise heterogeneity across wavelength bins is strongly tied to the instrument set-up. Thus, while the sensitivity and noise properties may be heterogeneous across the wavelength axis for the same spectrum, they are highly homogeneous across different spectra. Also considering the massive scale at which spectroscopic surveys are conducted, this makes the problem of spectral model fitting uniquely suited to be solved with ANPE.

The association of data properties with the observing spectrograph strongly indicates a program to develop ANPE models for the most popular spectrographs currently used, including APOGEE, LAMOST, GALAH, Gaia RVS, Keck/HIRES, LBT/PEPSI, Magellan/MIKE, etc., which will facilitate real-time analysis as soon as data is reduced. The training, distribution, and application of such amortized inference models require user-friendly frameworks for NPE. We have developed one such framework and open-source software named *nbi* (*neural bayesian inference*), which appears in a separate paper in this workshop.

To summarize, the *nbi* framework tackles two key issues that have greatly slowed the integration of NPE methods into our routine data analysis toolbox. The first issue dealt with *nbi* is the inexact nature of Neural Posterior Estimation, which we refer the reader to the other paper for details. The second aspect is the steep learning curve of NPE for the domain person less familiar for deep learning. To apply existing software frameworks supporting NPE such as *sbi* (Tejero-

Cantero et al., 2020) and *LAMPE*<sup>2</sup>, one has to first and foremost develop their own featurizer (a.k.a. embedding/compressor) network that compresses the high-dimensional observed data into a low-dimensional vector, which can be “read in” by the normalizing flow as the conditional. This need for customization contrasts sharply with the limited data types that astronomers regularly face, in particular sequential data such as light curves and spectra.

In our previous work of applying ANPE to binary-lens microlensing (Zhang et al., 2021), we have designed a hybrid CNN/RNN network (ResNet-GRU) for light-curve data, which is more effective than either ResNet or GRU alone. As demonstrated in the current work, this network architecture is also effective for spectral data. Therefore, the *nbi* package integrates this ResNet-GRU network architecture<sup>3</sup>, thus enabling off-the-shelf application to inference on sequential data. By applying the *nbi* package, we were able to develop an ANPE model for the APOGEE survey with minimal customization efforts.

This paper serves as a proof-of-concept for the benefits of developing spectral inference models tied to particular instruments using the ANPE method and the *nbi* package. We demonstrate the efficacy of our approach by training the *nbi* package on simulated spectra produced using *The Payne* (Ting et al., 2019), with realistic noise properties from the APOGEE survey. We demonstrate satisfactory results both on simulated APOGEE spectra and real APOGEE spectra. Our approach may be expanded to multiple instruments to serve as powerful tools for the spectroscopy community.

## 2. Data

The Apache Point Observatory Galactic Evolution Experiment (Majewski et al., 2017) is a near-infrared (NIR) survey with a spectral resolving power of  $R \sim 22,500$ . APOGEE observed stars in both the northern and southern hemispheres and its NIR spectra have been heavily used to study stellar evolution, stellar chemical abundances, and Galactic structure (Helmi et al., 2018; Price-Whelan et al., 2018). Here, we use data from the 17th data release. APOGEE DR17 contains  $\sim 657,000$  unique stars and  $\sim 2,660,000$  individual spectra (Abdurro’uf et al., 2022).

To model the APOGEE spectra, we adopt *The Payne* (Ting et al., 2019) which provides a simple multi-layer perceptron (MLP) model to interpolate between synthetic NIR spectra. *The Payne* yields spectral flux predictions with median absolute errors of  $\sim 10^{-3}$  for most test stars. *The Payne* model consists of 25 labels consisting of stellar parameters and elemental abundances, including  $T_{\text{eff}}$ ,  $\log(g)$ ,  $[\text{Fe}/\text{H}]$ ,

$v_{\text{micro}}$ ,  $v_{\text{macro}}$ ,  $C_{12}/C_{13}$  and  $[\text{X}/\text{Fe}]$  where  $\text{X} = \text{C}, \text{N}, \text{O}, \text{Na}, \text{Mg}, \text{Al}, \text{Si}, \text{P}, \text{S}, \text{K}, \text{Ca}, \text{Ti}, \text{V}, \text{Cr}, \text{Mn}, \text{Fe}, \text{Co}, \text{Ni}, \text{Cu}, \text{and Ge}$ . *The Payne* outputs continuum-normalized spectral fluxes for 7214 wavelength grid points.

We generate a training set of  $\sim 10^6$  spectra using *The Payne*. We sample 103,344 sets of  $T_{\text{eff}}$ ,  $\log(g)$  and  $C_{12}/C_{13}$  from the MIST isochrones (Choi et al., 2016) for stars with ages between 1–10 Gyr and where  $T_{\text{eff}} \in [3000 \text{ K}, 8000 \text{ K}]$ . Next, for each set of  $T_{\text{eff}}$ ,  $\log(g)$  and  $C_{12}/C_{13}$ , we randomly sample 10 instances of the remaining labels ( $v_{\text{micro}}$ ,  $v_{\text{macro}}$ ,  $[\text{Fe}/\text{H}]$ ,  $[\text{X}/\text{Fe}]$ ) using the priors from Ting et al. (2019). For the test set, we select 995 stars from APOGEE DR17 without any data reduction flags. We continuum-normalize the observed spectra and apply a mask to select pixels for fitting using the utilities provided by *The Payne*<sup>4</sup>. Each spectrum has 7214 wavelength pixels where each pixel is associated with a continuum-normalized flux, continuum-normalized flux error and a binary bitmask flag. The pixel bitmasks from the published APOGEE spectra can take a wide variety of values and is used to track problematic pixels in the observed spectra, but for simplicity, we convert this to a binary flag where GOOD=0 and BAD=1.

## 3. Methods and Results

This work exemplifies the simplicity of applying the *nbi* package for inference of sequential data in astronomy. We utilize the built-in “ResNet-GRU” network for dimensionality reduction, with the following hyperparameters: `dim_conv_max=512`, `depth=8`. We adapt the following parameters for the flow: `flow_hidden=256`, `num_cond_inputs=256`, `num_blocks=15`. After determining these values, the remaining required customization under *nbi* is writing a function describing the noise properties, which will be called during training to apply to the noise-free simulations on the fly.

Crucial to a successful NPE application is that the training data should resemble the test data—the real spectra taken by APOGEE. Therefore, we directly use the noise (error-bars) of the our 995 test-set APOGEE spectra to be applied during training. During training, we randomly re-scale each noise sequence such that the average S/N is uniformly drawn from 25 and 600—representative of the test set. To account for the aforementioned “bad-pixel” bit-masks, we also adapt the 995 bit-masks of our test-set data to apply during training. Specifically, we replace each masked value with the nearest non-masked value towards the left. If the first value is masked, we then replace it with the nearest non-masked value to the right. The particular inpainting scheme does not matter, so long as it is consistent, and the aforementioned scheme is very easy to implement in practice. We also

<sup>2</sup><https://github.com/francois-rozet/lampe>

<sup>3</sup>more network architectures will be provided in the future to facilitate applying *nbi* to other data types.

<sup>4</sup>[https://github.com/tingyuansen/The\\_Payne](https://github.com/tingyuansen/The_Payne)

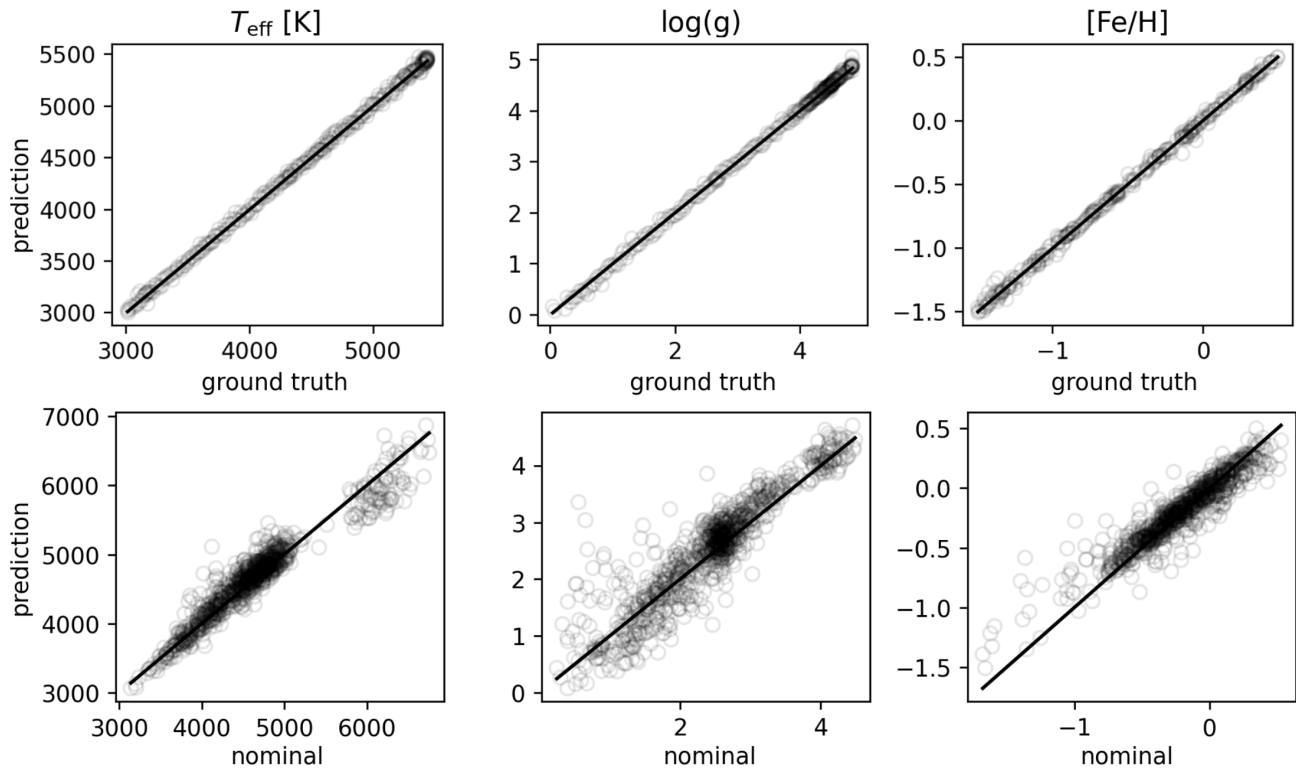


Figure 1. Top: Prediction vs. ground truth for a test set of simulated spectra with realistic APOGEE noise properties. Bottom: Prediction vs. nominal (DR 17) parameters retrieved from the APOGEE catalog, for a test set of 995 real APOGEE spectra. Note that these catalog parameters are not ground truths (see text).




Figure 2. Example fitting of a real APOGEE spectra using *nbi*. Red is the best-fit spectra model and black is observed spectra with error-bars. Masked values are not shown.

include the bit-mask as another input channel.

We then tested the trained model using both mock spectra unseen during training, as well as real APOGEE spectra. For both testing, we apply the same aforementioned masking and inpainting scheme. In Figure 1, we show test results for the three most important parameters— $T_{\text{eff}}$ ,  $\log(g)$  and  $[\text{Fe}/\text{H}]$ —out of the total of 26 parameters inferred upon. Unsurprisingly, performance on the mock test set is nearly perfectly aligned to the diagonal. Since NPE for well-specified models is a rather mature problem, we do not further discuss here, and instead focus on testing on real observed spectra.

In the lower three subplots in Figure 1, we apply the trained model to real APOGEE observed spectra and compare the *nbi* predictions to the uncalibrated DR17 labels obtained from the ASPCAP pipeline (García Pérez et al., 2016). First, we point out that the DR17 labels are not derived using the *The Payne* as the forward model. Therefore, they are not ground truths as far as this testing concerns, and the increased level of scattering is to a large extent attributed to the difference in the adopted forward model. To exemplify accuracy of the *nbi* derived parameters, we simply show an example of posterior-predictive (best-fit) spectra against the actual observed spectra (Figures 2&3). The *nbi* predictive spectra using *The Payne* accurately captured the observed

spectral features.

Moreover, we found that over the full test set of 995 real APOGEE spectra, the fitting residuals are smaller for 95% of the time for *nbi*-derived parameters compared to the DR 17 APOGEE parameters. Again, we caution the difference in the adapted forward models, but this provides preliminary evidence as to the effectiveness of the NPE method and the *nbi* framework for this task. We note that the  $[\text{Cu}/\text{H}]$  and  $[\text{Ge}/\text{H}]$  parameters were not provided in DR17, and we simply use the same *nbi* predicted values in the above testing. However, our testing on mock spectra reveals that the *nbi* model was able to provide meaningful constraint on these parameters. Therefore, the application of our trained *nbi* model will be able to produce a value-added catalog when applied to the full APOGEE survey.

#### 4. Discussion

In this paper, we first discuss the utility and appropriateness in training ANPE models for the problem of inference on spectral data. Due to the homogeneous noise properties for spectra taken with the instrument, ANPE is a natural fit in solving these problems. We introduce a training-time data-processing scheme to handle the noise properties, specifically by adapting the actual error-bars and bit-masks for




Figure 3. Same example fitting as seen in Figure 2, but zoomed in to show details.

observed spectra. The data-handling approach presented here is equally applicable to spectral fitting of galaxies, exoplanet atmospheres, etc. By applying our *nbi* framework concurrently submitted to this workshop, we were able to acquire the ANPE model with minimal customization efforts.

The current work serves as a proof-of-concept for a larger program to develop ANPE models for additional spectrographs, such as LAMOST, GALAH, *Gaia* RVS, Keck/HIRES, LBT/PEPSI, Magellan/MIKE, etc., which will greatly benefit analysis of such data. After the trained models are uploaded to the cloud, others may easily load and apply the model with *nbi* framework. The user may directly use the trained models to acquire accurate spectral model parameters in real-time, rather than waiting  $> 30$  minutes in the case of commonly used tools for spectral fitting like *iSpec* (Blanco-Cuaresma et al., 2014), providing a massive speedup.

### Acknowledgment

This work is supported by NSF award #2206744: “CDS&E: Accelerating Astrophysical Insight at Scale with Likelihood-Free Inference.” Support for TJ was provided by NASA through the NASA Hubble Fellowship grant HF2-51509

awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555.

### References

Abdurro’uf, Accetta, K., Aerts, C., Silva Aguirre, V., Ahumada, R., Ajaonkar, N., Filiz Ak, N., Alam, S., Allende Prieto, C., Almeida, A., Anders, F., Anderson, S. F., Andrews, B. H., Anguiano, B., Aquino-Ortíz, E., Aragón-Salamanca, A., Argudo-Fernández, M., Ata, M., Aubert, M., Avila-Reese, V., Badenes, C., Barbá, R. H., Barger, K., Barrera-Ballesteros, J. K., Beaton, R. L., Beers, T. C., Belfiore, F., Bender, C. F., Bernardi, M., Bershady, M. A., Beutler, F., Bidin, C. M., Bird, J. C., Bizyaev, D., Blanc, G. A., Blanton, M. R., Boardman, N. F., Bolton, A. S., Boquien, M., Borissova, J., Bovy, J., Brandt, W. N., Brown, J., Brownstein, J. R., Brusa, M., Buchner, J., Bundy, K., Burchett, J. N., Bureau, M., Burgasser, A., Cabang, T. K., Campbell, S., Cappellari, M., Carlberg, J. K., Wanderley, F. C., Carrera, R., Cash, J., Chen, Y.-P., Chen, W.-H., Cherinka, B., Chiappini, C., Choi, P. D., Chojnowski, S. D., Chung, H., Clerc, N., Cohen, R. E., Comerford, J. M., Comparat, J., da Costa, L., Covey, K., Crane, J. D., Cruz-Gonzalez, I., Culhane, C., Cunha, K., Dai, Y. S.,



- Damke, G., Darling, J., Davidson, James W., J., Davies, R., Dawson, K., De Lee, N., Diamond-Stanic, A. M., Cano-Díaz, M., Sánchez, H. D., Donor, J., Duckworth, C., Dwelly, T., Eisenstein, D. J., Elsworth, Y. P., Em-sellem, E., Eracleous, M., Escoffier, S., Fan, X., Farr, E., Feng, S., Fernández-Trincado, J. G., Feuillet, D., Filipp, A., Fillingham, S. P., Frinchaboy, P. M., Fromenteau, S., Galbany, L., García, R. A., García-Hernández, D. A., Ge, J., Geisler, D., Gelfand, J., Géron, T., Gibson, B. J., Goddy, J., Godoy-Rivera, D., Grabowski, K., Green, P. J., Greener, M., Grier, C. J., Griffith, E., Guo, H., Guy, J., Hadjara, M., Harding, P., Hasselquist, S., Hayes, C. R., Hearty, F., Hernández, J., Hill, L., Hogg, D. W., Holtzman, J. A., Horta, D., Hsieh, B.-C., Hsu, C.-H., Hsu, Y.-H., Huber, D., Huertas-Company, M., Hutchinson, B., Hwang, H. S., Ibarra-Medel, H. J., Chitham, J. I., Ilha, G. S., Imig, J., Jaekle, W., Jayasinghe, T., Ji, X., Johnson, J. A., Jones, A., Jönsson, H., Katkov, I., Khalatyan, Arman, D., Kinemuchi, K., Kisku, S., Knapen, J. H., Kneib, J.-P., Kollmeier, J. A., Kong, M., Kounkel, M., Kreckel, K., Krishnarao, D., Lacerna, I., Lane, R. R., Langgin, R., Lavender, R., Law, D. R., Lazarz, D., Leung, H. W., Leung, H.-H., Lewis, H. M., Li, C., Li, R., Lian, J., Liang, F.-H., Lin, L., Lin, Y.-T., Lin, S., Lintott, C., Long, D., Longa-Peña, P., López-Cobá, C., Lu, S., Lundgren, B. F., Luo, Y., Mackereth, J. T., de la Macorra, A., Mahadevan, S., Majewski, S. R., Manchado, A., Mandeville, T., Maraston, C., Margalef-Bentabol, B., Masseron, T., Masters, K. L., Mathur, S., McDermid, R. M., Mckay, M., Merloni, A., Merrifield, M., Meszaros, S., Miglio, A., Di Mille, F., Minniti, D., Minsley, R., Monachesi, A., Moon, J., Mosser, B., Mulchaey, J., Muna, D., Muñoz, R. R., Myers, A. D., Myers, N., Nadathur, S., Nair, P., Nandra, K., Neumann, J., Newman, J. A., Nidever, D. L., Nikakhtar, F., Nitschelm, C., O’Connell, J. E., Garma-Oehmichen, L., Luan Souza de Oliveira, G., Olney, R., Oravetz, D., Ortigoza-Urdaneta, M., Osorio, Y., Otter, J., Pace, Z. J., Padilla, N., Pan, K., Pan, H.-A., Parikh, T., Parker, J., Peirani, S., Peña Ramírez, K., Penny, S., Percival, W. J., Perez-Fournon, I., Pinsonneault, M., Poidevin, F., Poovelil, V. J., Price-Whelan, A. M., Bárbara de Andrade Queiroz, A., Rad-dick, M. J., Ray, A., Rembold, S. B., Riddle, N., Riffel, R. A., Riffel, R., Rix, H.-W., Robin, A. C., Rodríguez-Puebla, A., Roman-Lopes, A., Román-Zúñiga, C., Rose, B., Ross, A. J., Rossi, G., Rubin, K. H. R., Salvato, M., Sánchez, S. F., Sánchez-Gallego, J. R., Sanderson, R., Santana Rojas, F. A., Sarceno, E., Sarmiento, R., Sayres, C., Sazonova, E., Schaefer, A. L., Schiavon, R., Schlegel, D. J., Schneider, D. P., Schultheis, M., Schwöpe, A., Serenelli, A., Serna, J., Shao, Z., Shapiro, G., Sharma, A., Shen, Y., Shetrone, M., Shu, Y., Simon, J. D., Skrutskie, M. F., Smethurst, R., Smith, V., Sobek, J., Spoo, T., Sprague, D., Stark, D. V., Stassun, K. G., Steinmetz, M., Stello, D., Stone-Martinez, A., Storchi-Bergmann, T., Stringfellow, G. S., Stutz, A., Su, Y.-C., Taghizadeh-Popp, M., Talbot, M. S., Tayar, J., Telles, E., Teske, J., Thakar, A., Theissen, C., Tkachenko, A., Thomas, D., Tojeiro, R., Hernandez Toledo, H., Troup, N. W., Trump, J. R., Trussler, J., Turner, J., Tuttle, S., Unda-Sanzana, E., Vázquez-Mata, J. A., Valentini, M., Valenzuela, O., Vargas-González, J., Vargas-Magaña, M., Alfaro, P. V., Villanova, S., Vincenzo, F., Wake, D., Warfield, J. T., Washington, J. D., Weaver, B. A., Weijmans, A.-M., Weinberg, D. H., Weiss, A., Westfall, K. B., Wild, V., Wilde, M. C., Wilson, J. C., Wilson, R. F., Wilson, M., Wolf, J., Wood-Vasey, W. M., Yan, R., Zamora, O., Zaslowski, G., Zhang, K., Zhao, C., Zheng, Z., Zheng, Z., and Zhu, K. The Seventeenth Data Release of the Sloan Digital Sky Surveys: Complete Release of MaNGA, MaStar, and APOGEE-2 Data. , 259(2):35, April 2022. doi: 10.3847/1538-4365/ac4414.
- Blanco-Cuaresma, S., Soubiran, C., Heiter, U., and Jofré, P. Determining stellar atmospheric parameters and chemical abundances of FGK stars with iSpec. , 569:A111, September 2014. doi: 10.1051/0004-6361/201423945.
- Choi, J., Dotter, A., Conroy, C., Cantiello, M., Paxton, B., and Johnson, B. D. Mesa Isochrones and Stellar Tracks (MIST). I. Solar-scaled Models. , 823(2):102, June 2016. doi: 10.3847/0004-637X/823/2/102.
- Dax, M., Green, S. R., Gair, J., Macke, J. H., Buonanno, A., and Schölkopf, B. Real-Time Gravitational Wave Science with Neural Posterior Estimation. *Physical Review Letters*, 127(24):241103, December 2021. doi: 10.1103/PhysRevLett.127.241103. URL <https://link.aps.org/doi/10.1103/PhysRevLett.127.241103>. Publisher: American Physical Society.
- García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., Shetrone, M., Mészáros, S., Bizyaev, D., Carrera, R., Cunha, K., García-Hernández, D. A., Johnson, J. A., Majewski, S. R., Nidever, D. L., Schiavon, R. P., Shane, N., Smith, V. V., Sobek, J., Troup, N., Zamora, O., Weinberg, D. H., Bovy, J., Eisenstein, D. J., Feuillet, D., Frinchaboy, P. M., Hayden, M. R., Hearty, F. R., Nguyen, D. C., O’Connell, R. W., Pinsonneault, M. H., Wilson, J. C., and Zasowski, G. ASPCAP: The APOGEE Stellar Parameter and Chemical Abundances Pipeline. , 151(6): 144, June 2016. doi: 10.3847/0004-6256/151/6/144.
- Hahn, C. and Melchior, P. Accelerated Bayesian SED Modeling Using Amortized Neural Posterior Estimation. *The Astrophysical Journal*, 938(1):11, October 2022. ISSN 0004-637X. doi: 10.3847/1538-4357/ac7b84. URL <https://dx.doi.org/10.3847/>

- 1538-4357/ac7b84. Publisher: The American Astronomical Society.
- Helmi, A., Babusiaux, C., Koppelman, H. H., Massari, D., Veljanoski, J., and Brown, A. G. A. The merger that led to the formation of the Milky Way’s inner stellar halo and thick disk. , 563(7729):85–88, October 2018. doi: 10.1038/s41586-018-0625-x.
- Khullar, G., Nord, B., Ćiprijanović, A., Poh, J., and Xu, F. DIGS: deep inference of galaxy spectra with neural posterior estimation. *Machine Learning: Science and Technology*, 3(4):04LT04, December 2022. ISSN 2632-2153. doi: 10.1088/2632-2153/ac98f4. URL <https://dx.doi.org/10.1088/2632-2153/ac98f4>. Publisher: IOP Publishing.
- Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., Al-lende Prieto, C., Barkhouser, R., Bizyaev, D., Blank, B., Brunner, S., Burton, A., Carrera, R., Chojnowski, S. D., Cunha, K., Epstein, C., Fitzgerald, G., García Pérez, A. E., Hearty, F. R., Henderson, C., Holtzman, J. A., Johnson, J. A., Lam, C. R., Lawler, J. E., Maseman, P., Mészáros, S., Nelson, M., Nguyen, D. C., Nidever, D. L., Pinsonneault, M., Shetrone, M., Smee, S., Smith, V. V., Stolberg, T., Skrutskie, M. F., Walker, E., Wilson, J. C., Zasowski, G., Anders, F., Basu, S., Beland, S., Blanton, M. R., Bovy, J., Brownstein, J. R., Carlberg, J., Chaplin, W., Chiappini, C., Eisenstein, D. J., Elsworth, Y., Feuillet, D., Fleming, S. W., Galbraith-Frew, J., García, R. A., García-Hernández, D. A., Gillespie, B. A., Girardi, L., Gunn, J. E., Hasselquist, S., Hayden, M. R., Hekker, S., Ivans, I., Kinemuchi, K., Klaene, M., Mahadevan, S., Mathur, S., Mosser, B., Muna, D., Munn, J. A., Nichol, R. C., O’Connell, R. W., Parejko, J. K., Robin, A. C., Rocha-Pinto, H., Schultheis, M., Serenelli, A. M., Shane, N., Silva Aguirre, V., Sobek, J. S., Thompson, B., Troup, N. W., Weinberg, D. H., and Zamora, O. The Apache Point Observatory Galactic Evolution Experiment (APOGEE). , 154(3):94, September 2017. doi: 10.3847/1538-3881/aa784d.
- Price-Whelan, A. M., Hogg, D. W., Rix, H.-W., De Lee, N., Majewski, S. R., Nidever, D. L., Troup, N., Fernández-Trincado, J. G., García-Hernández, D. A., Longa-Peña, P., Nitschelm, C., Sobek, J., and Zamora, O. Binary Companions of Evolved Stars in APOGEE DR14: Search Method and Catalog of ~5000 Companions. , 156(1):18, July 2018. doi: 10.3847/1538-3881/aac387.
- Tejero-Cantero, A., Boelts, J., Deistler, M., Lueckmann, J.-M., Durkan, C., Gonçalves, P. J., Greenberg, D. S., and Macke, J. H. sbi: A toolkit for simulation-based inference. *Journal of Open Source Software*, 5(52):2505, August 2020. ISSN 2475-9066. doi: 10.21105/joss.02505. URL <https://joss.theoj.org/papers/10.21105/joss.02505>.
- Ting, Y.-S., Conroy, C., Rix, H.-W., and Cargile, P. The Payne: Self-consistent ab initio Fitting of Stellar Spectra. , 879(2):69, July 2019. doi: 10.3847/1538-4357/ab2331.
- Vasist, M., Rozet, F., Absil, O., Mollière, P., Nasedkin, E., and Loupe, G. Neural posterior estimation for exoplanetary atmospheric retrieval. *Astronomy & Astrophysics*, 672:A147, April 2023. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202245263. URL <https://www.aanda.org/articles/aa/abs/2023/04/aa45263-22/aa45263-22.html>. Publisher: EDP Sciences.
- Zhang, K., Bloom, J. S., Gaudi, B. S., Lanusse, F., Lam, C., and Lu, J. R. Real-time Likelihood-free Inference of Roman Binary Microlensing Events with Amortized Neural Posterior Estimation. *The Astronomical Journal*, 161(6):262, May 2021. ISSN 1538-3881. doi: 10.3847/1538-3881/abf42e. URL <https://dx.doi.org/10.3847/1538-3881/abf42e>. Publisher: The American Astronomical Society.